PPCEGMFinal.rtf
7/18/10

Collective Epistemology, eds. Hans Bernard Schmid, Marcel Weber, and Daniel Sirtes, Ontos Verlag, (2011): 157-175.

Probabilistic Proofs and the Collective Epistemic Goals of Mathematicians

Don Fallis
School of Information Resources
University of Arizona
fallis@email.arizona.edu

Abstract

Mathematicians only use *deductive proofs* to establish that mathematical claims are true. They never use *inductive evidence*, such as *probabilistic proofs*, for this task. Don Fallis (1997 and 2002) has argued that mathematicians do not have good epistemic grounds for this complete rejection of probabilistic proofs. But Kenny Easwaran (2009) points out that there is a gap in this argument. Fallis only considered how mathematical proofs serve the epistemic goals of *individual* mathematicians. Easwaran suggests that deductive proofs might be epistemically superior to probabilistic proofs because they are *transferable*. That is, one mathematician can give such a proof to another mathematician who can then verify for herself that the mathematical claim in question is true without having to rely at all on the testimony of the first mathematician. In this paper, I argue that *collective epistemic goals* are critical to understanding the methodological choices of mathematicians. But I argue that the collective epistemic goals promoted by transferability do not explain the complete rejection of probabilistic proofs.

Introduction

In order to explain why people engage in the activities that they do, we often attribute goals and desires to them. For example, we hypothesize that a student wants to pass an exam in order to explain why she is studying so hard. And this technique can be used to explain the behavior of groups as well as individuals. For example, philosophers (e.g., Maddy 1997, Goldman 1999, 221-71) have tried to identify the goals of scientists and mathematicians that will explain their methodological choices.[1]

In this paper, I use this technique to try to understand one particular methodological choice of mathematicians. Namely, mathematicians will only use deductive proofs to establish that mathematical claims are true. Mathematicians will not use inductive evidence, such as *probabilistic proofs*, for this task (cf. Detlefsen and Luker 1980, 818-19, Fallis 2002, 374-76, Peressini 2003). The project here is to see if we can find a goal that will explain the complete rejection of probabilistic proofs.

---

[1] See Fallis (2002, 384-85) for a defense of using this technique to explain the methodological choices of mathematicians.

Understanding why mathematicians only use deductive proofs to establish that mathematical claims are true is clearly critical for the philosophy of mathematics. But it is also an important issue in the philosophy of science more generally and even in epistemology. Philosophers (e.g., Descartes) as well as mathematicians clearly think that deductive evidence is epistemically superior to inductive evidence (cf. Couvalis 2004, 28-29). But this is largely an unexamined article of faith. Thus, it would be useful to explicitly identify the goals that deductive evidence, but not inductive evidence, allows us to achieve.

But before we get started, there are a few important points about this project that should be clarified at the outset. First, the methodological choices of scientists and mathematicians can potentially be explained by pragmatic goals as well as by epistemic goals. For instance, it may be that mathematicians who use probabilistic proofs cannot get their papers published or get funding from the NSF. However, following most other philosophers of science and mathematics, the project here is to identify *epistemic* goals (e.g., the goals of acquiring true beliefs, avoiding false beliefs, gaining understanding) that explain the complete rejection of probabilistic proofs.

Second, there are actually many epistemic goals that deductive proofs *often* serve, but that probabilistic proofs *almost never* serve (e.g., as I discuss below, the goal of gaining understanding). However, a deductive proof can *always* be used to establish that a mathematical claim is true. Thus, in order to explain the *complete* rejection of probabilistic proofs, we need to find an epistemic goal that deductive proofs *always* serve, but that probabilistic proofs *never* serve.

Third, methodological choices can potentially be explained by the goals that individuals have in virtue of being members of a scientific community or by the goals that the scientific community itself has. For example, just as individual mathematicians only want to accept mathematical claims that are true, the mathematical community as a whole only wants to accept mathematical claims that are true. In trying to explain the complete rejection of probabilistic proofs, I will consider both types of *collective epistemic goals* (cf. Fallis 2007).

Finally, it should be noted that mathematicians often do make use of inductive evidence (cf. Corfield 2003, 103-29, Avigad 2008, 308-12, Bledin 2008, 497-98, Baker 2009). For example, it is an important part of why mathematicians believe that certain axioms are true (cf. Whitehead and Russell 1962 [1910], 59, Maddy 1997, 36-62). Also, it can help to identify plausible conjectures that it might be worth trying to prove (cf. Brown 1999, 158-71). But the project here is to explain why inductive evidence is never used in place of deductive proofs to establish that mathematical claims are true.[2]

---

[2] Mathematicians do not always prove mathematical claims in order to establish that they are true (cf. Dawson 2006, 275-81). For example, mathematicians commonly search for simpler and more elegant proofs of results that have already been established. In fact, they prove results that are much more certain than the axioms from which they are derived, as when Whitehead and Russell (1962 [1910]) took several hundred pages to formally prove that 1 plus 1 equals 2. But the issue here is whether deductive proofs are superior with respect to establishing that mathematical claims are true.

Two Proofs of Primality

For the sake of concreteness, I will focus my discussion on two specific methods that mathematicians might use to establish that some number *n* is prime. A very simple (although time consuming) way to prove *deductively* that *n* is prime is to try dividing *n* by each of the numbers less than *n*. If none of these numbers divide into *n* without a remainder, we conclude that *n is* prime. Mathematicians are perfectly happy to use this *trial division test* to establish that a number is prime.

A very efficient way to prove *probabilistically* that *n* is prime has been suggested by Michael Rabin (1980). Rabin proved (deductively) that, if a number *n* is *not* prime, then most (over 75%) of the numbers less than *n* are "witnesses" to this fact. That is, if *n* is not prime, then most smaller numbers have a special property *X* that can be checked for very quickly and that implies that *n* is not prime.[3] So, if we want to establish that *n* is prime, we pick a whole bunch of numbers less than *n* at random and check to see if any of them have property *X*. If none of these numbers "testify" that *n* is not prime, we conclude that *n is* prime. Unlike the trial division test, the *Rabin test* has to make use of a randomization device (such a coin flip) to determine exactly which calculations to perform. Although the Rabin test is extremely reliable, mathematicians will not use it to establish that a number is prime.

Explaining the Complete Rejection of Probabilistic Proofs

Several people (e.g., Peressini 2003, Avigad 2008) have pointed out that there is a clear distinction between deductive proofs and probabilistic proofs. *In principle*, a deductive proof provides a *guarantee* that a mathematical claim is true. That is, if no mistakes are made in the construction of a deductive proof, then there must be a valid derivation of the mathematical claim in question. By contrast, even if all of the calculations in a probabilistic proof are performed correctly, the mathematical claim in question might turn out to be false (e.g., because we got unlucky and missed all of the witnesses when we picked numbers at random).

However, mathematicians live in the real world of fallibility rather than the ideal world of infallibility (cf. Devlin 2004). As David Hume (1967 [1739], book I, part IV, section I) famously noted, "in all demonstrative sciences the rules are certain and infallible; but when we apply them, our fallible and uncertain faculties are very apt to depart from them, and fall into error." In fact, there are plenty of examples where mathematicians falsely believed (for quite a long time) that they had found a correct deductive proof of an important mathematical result (cf. Brown 1999, 156). Thus, deductive proofs in general do not actually provide mathematicians with certainty that mathematical claims are true.[4]

---

[3] See Rabin (1980, 130) for a precise mathematical statement of what property *X* is.

[4] Interestingly, George Couvalis (2004) argues that we can only know that we have deductively proved a mathematical claim by appealing to inductive evidence. In particular, our confidence that we have not made a mistake in this instance is based on the fact that we have rarely made mistakes in the past.

In order to explain the methodological choice of mathematicians, we need to find an *actual* epistemic benefit that deductive proofs provide and that probabilistic proofs do not (cf. Fallis 2002, 383).[5] However, I have argued previously (in Fallis 1997 and Fallis 2002) that there are no such epistemic benefits. For example, while deductive proofs may not provide absolute certainty that mathematical claims are true, it might be suggested that probabilistic proofs do not provide us with *enough* certainty. However, the Rabin test can provide as much certainty as might be required (cf. Fallis 2002, 380, Easwaran 2009). We just need to pick enough numbers at random and check to see if any of them have property *X*. Moreover, the Rabin test can provide more certainty than many deductive proofs, such as the trial division test. Because the trial division test requires more calculations than the Rabin test, there is a greater chance of calculation errors and, thus, a greater chance overall that *n* is not actually prime even if the test says that it is.

It might also be suggested that probabilistic proofs do not allow us to *understand* why a mathematical claim is true. Understanding is certainly an important epistemic value in mathematics (cf. Avigad 2008, 312-14, Bledin 2008, 498). We do not just want to know *that* something is true; we also want to know *why* it is true (cf. Wolfram 2002, 1156, Baker 2009, section 3.4). And, unlike probabilistic proofs, deductive proofs often do provide such understanding. However, whether we use a deductive proof or a probabilistic proof, there is not much understanding to be had when proving that a number is prime. With both methods, we simply understand that the number cannot be divided evenly by any smaller number.

In Fallis 1997 and Fallis 2002, I consider several further epistemic goals that mathematicians have. In each case, I argue that probabilistic proofs promote these goals as well as deductive proofs in at least some circumstances (e.g., when proving primality). Thus, none of these epistemic goals explain the complete rejection of probabilistic proofs. However, we do not have an exhaustive list of the epistemic goals of mathematicians to work with.[6] Thus, my argument that there is no important epistemic distinction between deductive proofs and probabilistic proofs is not conclusive.

---

[5] It has been suggested that mathematics is *by definition* a deductive science. As Jeremy Avigad (2008, 307) puts it, "inductive evidence is not the right sort of thing to provide mathematical knowledge, as it is commonly understood." If this is correct, then someone who used a probabilistic proof to establish the truth of a mathematical claim would not acquire *mathematical* knowledge. But even if this is correct, it still begs the question of whether there is a good epistemic reason to define mathematics in this way. Indeed, some mathematicians (e.g., Zeilberger 1993, Wolfram 2002, 792-95) claim that this is not a good way to define mathematics. It might also be suggested that *deductive* proofs are simply intrinsically valuable to mathematicians. If this is correct, then deductivity itself is an actual epistemic benefit that deductive proofs provide and that probabilistic proofs do not. However, deductive proof being an end in itself is a somewhat unsatisfying explanation of the complete rejection of probabilistic proofs. We might have hoped that mathematicians only use deductive proof because it is the most effective means to achieving some further epistemic goal (such as avoiding errors or finding errors that have already been made). Along similar lines, Alvin Goldman (1999, 78) has criticized Helen Longino for assigning "fundamental epistemic value to impartiality and nonarbitrariness" in science when we would have thought that they were valuable only because they "foster accuracy and truth."

[6] W. S. Anglin (1997, 85-127) does provide a fairly comprehensive list of the things that mathematicians value.

Easwaran on Transferability

Kenny Easwaran (2009) identifies an important gap in my argument. Like most other philosopher of mathematics, I focused exclusively on the epistemic goals of mathematicians as *solitary inquirers*. However, as Easwaran points out, there is an important *social* element to scientific and mathematical practice (cf. Kitcher 1993, 303-89, Goldman 1999, 221-71).[7] In particular, collective epistemic goals are arguably critical to understanding the methodological choices of mathematicians. Indeed, such a goal might explain the complete rejection of probabilistic proofs.

By paying attention to the social element of mathematical practice, Easwaran thinks that he has been able to identify the critical distinction between deductive proofs and probabilistic proofs. Easwaran points out that deductive proofs, but not probabilistic proofs, are *transferable*. A proof is transferable if one mathematician can give the proof to another mathematician who can then verify for herself that the mathematical claim in question is true without having to rely at all on the testimony of the first mathematician. In other words, a transferable proof "needs nothing outside itself to be convincing" (Tymoczko 1979, 59).

It is clear that deductive proofs are transferable. In order to convince herself that the mathematical claim in question is true without any reliance on the testimony, the receiving mathematician just has to check that each step in the proof is valid.[8] For example, in the case of the trial division test, the receiving mathematician simply has to check that each calculation has been performed correctly to convince herself that *n* is prime.

Probabilistic proofs, however, are not transferable. For example, in the case of the Rabin test, even if the receiving mathematician checks that each calculation has been performed correctly, she might still worry that *n* is not prime. The Rabin test is only convincing evidence that *n* is prime if the receiving mathematician knows that the smaller numbers checked for property *X* were chosen at random. Now, any particular set of numbers *might* have been picked at random. However, the set of numbers might instead have been *deliberately* chosen from those few numbers that lack property *X* when *n* is not prime. And it is difficult to see how the receiving mathematician can know that this did not happen without relying on the word of the original mathematician.

The original mathematician could certainly establish that the numbers checked for property *X* were chosen as random in much the same way that other scientists establish that their data has not been faked (cf. Fallis 2000, 269). In fact, Easwaran considers several ways in which the original mathematician might do so. But he concludes (correctly I think) that there is always a possibility of deception. In order to convince yourself that the numbers were picked at random, you ultimately have to rely on the

[7] Heretofore, philosophers of science have paid much more attention to this social element than have philosophers of mathematics.
[8] In order to do this, the receiving mathematician must have sufficient mathematical expertise.

testimony of the original mathematician (or on the testimony of someone else who can vouch for the original mathematician).

Admittedly, you could avoid relying on testimony by picking your own numbers at random and checking to see if any of them have property *X*. This would provide you with excellent, and completely non-testimonial, evidence that *n* is prime. But you would (almost certainly) not end up choosing exactly the same numbers as the original mathematician. Thus, your belief that *n* is prime would not be justified by exactly the same proof that justified the original mathematician's belief.

Thus, with transferability, Easwaran has identified a distinction between deductive proofs and probabilistic proofs. But we could already draw a distinction between these two types of proofs. For example, probabilistic proofs make use of a randomization device whereas deductive proofs do not. What we want is to identify an *epistemically important* distinction. That is, we want to find a property of deductive proofs that always allows such proofs to promote an epistemic goal of mathematicians better than probabilistic proofs. In particular, this is what Easwaran needs in order to "defend the practice" of rejecting of probabilistic proofs. In the remainder of this paper, I argue that transferability is not such a property.[9]

The transferability of a proof allows a mathematician to check the proof for herself rather than having to rely on the testimony of another mathematician. Thus, the obvious suggestion is that transferability is valuable because it allows an individual to be *epistemically autonomous*. In general terms, someone is epistemically autonomous if she believes on the basis of reasons that she herself has considered (cf. Fricker 2006). Claims of the value of epistemic autonomy go back to Descartes and Locke (cf. Fricker 2006, 225). In fact, standard advice in critical thinking courses is not to rely on authority (cf. Huemer 2005, 523).

There are at least two ways in which the epistemic autonomy provided by transferable proofs can be epistemically beneficial to mathematicians. First, it can be beneficial to the individual mathematician who is epistemically autonomous. (Transferability is defined in terms of more than one individual. But it might be that this property is primarily valuable to the individual who receives the transferable proof from another mathematician.) Second, it can be beneficial to the mathematical community whose members are epistemically autonomous. In the following two sections, I discuss these two ways in which transferable proofs can promote the collective epistemic goals of mathematicians. However, I argue that the same benefits can be achieved with probabilistic proofs. Thus, transferability does not explain the complete rejection of probabilistic proofs.

The Value of Transferability for Individuals

[9] Jeffrey C. Jackson (2009) also argues that Easwaran has failed to identify an epistemically important distinction between deductive proofs and probabilistic proofs.

An individual who is epistemically autonomous is potentially less prone to error (cf. Fricker 2006, 242). If a mathematician asserts that she has proven that a particular mathematical claim is true and you just have to take her word for it, you are open to several potential sources of error. For example, this mathematician may have made a mistake in carrying out the proof. In addition, this mathematician may even be trying to deceive you (e.g., in order to get unwarranted credit for having proven this claim). For example, in the case of the Rabin test, no matter how many numbers she claims to have chosen at random and checked to see if they have property *X*, you have to worry about honest mistakes and deception. However, if a mathematician provides you with a transferable proof of a claim, you can reduce (if not entirely eliminate) such worries by checking the proof for yourself (assuming that you have sufficient mathematical expertise to do so).

As several philosophers (e.g., Huemer 2005, Fricker 2006) have pointed out, while it may sound good in the abstract, epistemic autonomy in general is not all that it is cracked up to be. If you had to consider the (non-testimonial) reasons for every proposition before believing it, you would not end up believing very much. We just do not have the time and expertise to do this for every proposition that we need to know (cf. Hardwig 1991, 693-94). In addition, unless you happen to have a great deal of expertise in a particular area, you are probably better off (in terms of reducing the probability of error) consulting experts and relying on their testimony (cf. Mathiesen 2006, 143).[10]

But if you restrict your epistemic autonomy to a limited domain, such as an area of mathematics in which you have great expertise, it may not be such a bad thing. You will not be epistemically impoverished in general. And you will not be epistemically impoverished in this area of mathematics because you have sufficient expertise to check the proofs in this area for yourself. You might miss out on a few true beliefs that you could have acquired through the testimony of other mathematicians in this area (e.g., because you do not have time to check all of the proofs in this area for yourself). But this may not be such a great epistemic cost since mathematicians tend to be extremely epistemically risk averse. Mathematicians are happy to forgo additional knowledge if that is what it takes to avoid falling into error. In other words, they are willing to trade off what Alvin Goldman (1987) calls *power* for greater *reliability*.

Even so, the benefits of epistemic autonomy for individual mathematicians do not explain the complete rejection of probabilistic proofs. First of all, an individual does not have to be epistemic autonomous in order to have mathematical knowledge. In particular, a mathematician can know that a mathematical claim is true even if she has not surveyed the entire proof for herself. There are deductive proofs that are so long that no individual mathematician has surveyed, or could survey, the whole thing. The standard example is the "10,000 page proof" of the classification of finite simple groups (cf. Brown 1999,

[10] In general, people certainly value autonomy. But it is not clear that it is always rational for them to do so. For instance, people often prefer to drive themselves rather than fly even though it puts them at greater risk of injury and death (cf. Myers 2001). Similarly, it is not clear that epistemic autonomy is rational if it puts you at greater epistemic risk.

158).[11] In such cases, the justification that a mathematician has for believing that the mathematical claim in question is true is partially based on the testimony of other mathematicians (cf. Hardwig 1991, 695-96).[12] In particular, she has to trust that these other mathematicians correctly verified the parts of the proof that they were responsible for.

All that seems to be required for mathematical knowledge is that, for each piece of a proof, *somebody* has surveyed it (cf. Azzouni 1994, 166). In fact, the somebody who has surveyed some of the pieces might even be a computer. For example, in the case of the Four-Color Theorem, mathematicians used a computer to check the proof because there were too many cases for humans to check by hand (cf. Tymoczko 1979). Relying on a computer in this way to check a deductive proof is essentially like relying on testimony (cf. McEvoy 2008, 383).[13]

Another reason why the benefits of epistemic autonomy for individual mathematicians do not explain the complete rejection of probabilistic proofs is that probabilistic proofs actually allow mathematicians to be epistemically autonomous. While probabilistic proofs are not transferable, they are *reproducible*. For example, if someone claims to have proven that n is prime using the Rabin test, you can avoid relying on testimony by performing the Rabin test again for yourself. That is, you can pick your own numbers at random and check to see if any of them have property X. Although it will (almost certainly) not be exactly the same evidence that the original mathematician had, you will have excellent, and completely non-testimonial, evidence that n is prime. Admittedly, having to perform the Rabin test for yourself will require some effort on your part. But even if you are given a transferable proof of a mathematical claim, checking this proof is going to take some effort.[14,15]

---

[11] Admittedly, in principle, with unlimited time and expertise, one person could survey the entire proof of the classification of finite simple groups. But, as noted above, our concern here is with the *actual* practices of mathematicians and the epistemic goals that these practices promote.

[12] In mathematics, there is a huge amount of collaboration in terms of discovering mathematical results (cf. Fallis 2006, 204). But there is also sometimes collaboration in terms of justifying mathematical results.

[13] Jody Azzouni (1994, 169-71) argues, to the contrary, that, unlike testimony, reliance on a computer involves an appeal to empirical science. But see McEvoy (2008, 383-86).

[14] In addition, some steps will typically be left out of a published proof and the reader will have to figure out how to fill them in (cf. Fallis 2003).

[15] Epistemic autonomy can also allow mathematicians to be more *sensitive* to undercutting evidence. For example, if you know that the proof of theorem A relies on lemma B and you discover that lemma B is actually false, you will no longer be confident that theorem A is true. And such sensitivity may have epistemic value beyond simply enhancing your reliability. As Elizabeth Fricker (2006, 242) puts it, "sensitivity to defeating evidence … is usually taken to be the hallmark of belief which amounts to knowledge." However, since mathematicians do not have to be epistemically autonomous in order to have mathematical knowledge, they do not have to be sensitive to all possible defeating evidence either. In fact, even if a mathematician does not rely on the testimony of other mathematicians, she will not be sensitive to certain sorts of potential defeaters. For example, mathematicians rely on pencils and chalkboards (as well as their own memory) without having a complete understanding of how these mechanisms work (cf. Azzouni 1994, 160). In addition, if you perform the Rabin test for yourself, you will be sensitive to many sorts of potential defeaters. For example, you will know to be worried if your set of randomly chosen numbers has some unexpected property (e.g., if you happened to pick 1, 2, 3, …, x -1, x). Admittedly, any mechanism for randomly choosing numbers may leave us insensitive to certain sorts of potential defeaters.

The Value of Transferability for the Group

It often makes sense to attribute beliefs and knowledge to groups, such as the Supreme Court, the FBI, or the scientific community (cf. Hakli 2006). In particular, we frequently talk about whether particular mathematical claims (e.g., the Four-Color Theorem, Fermat's Last Theorem, the Poincaré Conjecture) are known by the mathematical community to be true. In addition, just as individual mathematicians do not want to accept mathematical claims that are false, the mathematical community as a whole does not want to accept false claims. In fact, philosophers of science (e.g., Kitcher 1993, Goldman 1999, 221-71) have tried to identify practices that allow the scientific community as a whole to accept fewer false claims (or more true claims). In particular, using transferable proofs arguably allows the mathematical community to accept fewer false claims.

If individual mathematicians are epistemically autonomous, the mathematical community as a whole is less prone to error. If a mathematician asserts that she has proven that a particular mathematical claim is true and everyone just has to take her word for it, the mathematical community has to worry about honest mistakes and deception. However, if a mathematician provides the community with a transferable proof of the claim, such worries can be reduced (if not entirely eliminated) by having several other mathematicians check the proof.

Since mathematicians are fallible, we do not want to have proofs checked by just one individual. Mathematics would be too insecure. As David Hume (1967 [1739], book I, part IV, section I) put it, "there is no Algebraist nor Mathematician so expert in his science, as to place entire confidence in any truth immediately upon his discovery of it, or regard it as any thing, but a mere probability. Every time he runs over his proofs, his confidence encreases; but still more by the approbation of his friends; and is rais'd to its utmost perfection by the universal assent and applauses of the, learned world." In addition, this sort of practice does not have the aforementioned drawbacks that epistemic autonomy can have for individuals. For example, the mathematical community usually has the *collective* time and expertise to do this for any important mathematical claim.

However, the goal of the mathematical community to avoid error does not explain the complete rejection of probabilistic proofs. For example, if someone claims to have proven that n is prime using the Rabin test, worries about honest mistakes and deception can be reduced by having several other mathematicians perform the test for themselves. These other mathematicians will (almost certainly) not end up picking the same numbers at random as the original mathematician. But this does not adversely affect the reliability of the practice.

Indeed, this is essentially the same sort of practice that takes place in science in general. The scientific community as a whole is more reliable because scientists are able to

---

But mathematicians were presumably not worried about probabilistic proofs because they were unsure of their ability to pick numbers at random.

replicate experiments in order to confirm scientific results.[16,17] And it is more reliable even though scientists do not replicate experiments using the very same random sample of atoms, subjects, etc. that the original scientist used. Scientists take a new random sample and then perform the experiment again.[18]

But just as with individuals, simply being right that a claim is true is not the only important concern for groups which have epistemic goals. In particular, such groups typically want to have a shared understanding of *why* that claim is true. For example, it is problematic if the scientific experts on a particular topic (e.g., climate change) agree that a claim is true, but are not able to agree on why that claim is true.[19] But if a mathematician has to provide the mathematical community with a transferable proof of a claim, other mathematicians can easily acquire the very same reasons for believing that the claim is true.[20]

Admittedly, mathematicians (even experts in the same area) do not always take advantage of transferability to check a proof for themselves.[21] Thus, transferability only insures that they *would* have the same reasons if they did check. In other words, even if different mathematicians have different (often testimonial) reasons for believing that a claim is true, transferability insures that the non-testimonial reasons that ground their belief is the same.

However, the goal of the mathematical community to have shared reasons (even of this counterfactual sort) does not explain the complete rejection of probabilistic proofs. If several mathematicians perform the Rabin test for themselves, their belief that *n* is prime will not be justified by exactly the same (non-testimonial) evidence. In particular, they

---

[16] As Philip Kitcher (1993, 336) puts it, "attempts at replication are frequently (though not always) a good thing for the *community*." Kitcher even considers exactly how many scientists should try to replicate experiments in order to confirm new results. His model does this by factoring in the initial probability that the new result is correct, the benefit of building on the new result if it is correct, the cost of building on the new result if it is incorrect, the cost of the time spent trying to replicate the new result, and the reliability of the scientists who are trying to replicate the new result.

[17] Actually having someone else replicate an experiment can enhance reliability by catching mistakes and deception. But in addition, the mere possibility of someone else replicating an experiment can enhance reliability by deterring negligence and deception in the first place (cf. Hardwig 1991, 706). Admittedly, it is not always feasible to replicate scientific experiments (cf. Hardwig 1991, 705-06). Even so, important results which other results build upon are tested indirectly. And this can also occur in mathematics regardless of whether a result was established using a deductive proof or a probabilistic proof.

[18] Probabilistic proofs in mathematics and (at least some forms of) induction in science both involve inferring that a full population has certain properties based on the fact that a much smaller, randomly chosen, sample of that population has those properties. But it should be noted that there are also important differences between the two techniques. For example, scientific induction typically involves inferring that the future will be like the past in certain ways, whereas any mathematical argument is presumably timeless.

[19] Similarly, when a court with several members (e.g., the Supreme Court) agrees on a decision, but not on why that decision is the correct one, the decision does not have as much legal force (cf. Warnken and Samuels 1997).

[20] As proofs are getting longer and more complicated, there can be disputes that are not all that easy to resolve (cf. De Millo et al. 1979, 272, Devlin 2004). As a result, using a deductive proof does not *guarantee* that mathematicians will end up with the very same reasons.

[21] As the example of the 10,000 page proof shows, it is not always even possible for them to do so.

each know, of a *different* set of numbers, that those numbers do not have property *X*. However, what ultimately justifies their belief that *n* is prime is that several numbers have been chosen *at random* and those numbers (whatever they happen to be) do not have property *X*. Thus, they have the same reasons for believing that *n* is prime in the very same sense that different scientists who have replicated an experiment have the same non-testimonial reasons for believing a scientific result.

While Easwaran is certainly correct that different mathematicians who independently perform the Rabin test will not have *exactly* the same evidence, it is not clear why this is epistemically important.[22] In addition, it is not even clear that transferable proofs always guarantee that different mathematicians will end up being justified by exactly the same evidence. According to Easwaran, *proof sketches* count as transferable proofs. In other words, a rough outline of a proof can give a mathematician enough information to verify for herself that a claim is true without relying on the testimony of the mathematician who came up with the proof. However, it is acceptable for mathematicians to leave very large gaps in such proof sketches (cf. Fallis 2003). Thus, there is no guarantee that different mathematicians, when they check the proof, will fill in these gaps in exactly the same way.

Conclusion

Easwaran has identified a distinction between deductive proofs and probabilistic proofs. Namely, deductive proofs are transferable, but probabilistic proofs are not. However, this distinction does not explain the complete rejection of probabilistic proofs. It is not the case that transferable proofs always promote a collective epistemic goal of mathematicians better than non-transferable proofs. That being said, it is important to consider the collective epistemic goals of mathematicians when trying to explain their methodological choices. And it may still turn out to be the case that such a goal will ultimately provide an explanation for the complete rejection of probabilistic proofs.[23]

References

Anglin, W. S. 1997. *The Philosophy of Mathematics*. Lewiston, New York: Edwin Mellen Press.

Avigad, Jeremy. 2008. "Computers in Mathematical Inquiry." Pp. 302-16 in *The Philosophy of Mathematical Practice*, ed. Paolo Mancosu. Oxford: Oxford.

[22] It is actually rather difficult to say when two proofs are exactly the same (cf. Dawson 2006, 272-75). It may be necessary to do so in order to settle priority disputes, such as determining who deserves credit for proving the Poincaré Conjecture (cf. Nasar and Gruber 2006). But settling priority disputes does not improve our *mathematical* knowledge.

[23] I would especially like to thank Kenny Easwaran and Kay Mathiesen for many helpful discussions on this topic. I would also like to thank Jeff Jackson for his comments on a draft of this paper. Finally, I would like to thank the organizers (Hans Bernard Schmid, Marcel Weber, and Daniel Sirtes) and the participants at the *Workshop on Collective Epistemology*, Universität Basel, 2008 for their feedback.

Azzouni, Jody. 1994. *Metaphysical Myths, Mathematical Practice*. New York: Cambridge.

Baker, Alan. 2009. "Non-Deductive Methods in Mathematics." *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/mathematics-nondeductive/

Bledin, Justin. 2008. "Challenging Epistemology: Interactive Proofs and Zero Knowledge." *Journal of Applied Logic* 6:490-501.

Brown, James R. 1999. *Philosophy of Mathematics*. New York: Routledge.

Corfield, David. 2003. *Towards a Philosophy of Real Mathematics*. Cambridge: Cambridge.

Couvalis, George. 2004. "Is Induction Epistemologically Prior to Deduction?" *Ratio* 17:28-44.

Dawson, John W. 2006. "Why Do Mathematicians Re-Prove Theorems?" *Philosophia Mathematica* 14:269-86.

De Millo, Richard, Richard Lipton, and Alan Perlis. 1979. "Social Processes and Proofs of Theorems and Programs." *Communications of the Association for Computing Machinery* 22:271-80.

Detlefsen, Michael and Mark Luker. 1980. "The Four-Color Theorem and Mathematical Proof." *Journal of Philosophy* 76:803-20.

Devlin, Keith. 2004. "Mathematicians Face Uncertainty." *Discover* 25(1):36.

Easwaran, Kenny. 2009. "Probabilistic Proofs and Transferability." *Philosophia Mathematica* 17:341-62.

Fallis, Don. 1997. "The Epistemic Status of Probabilistic Proof." *Journal of Philosophy* 94:165-86.

Fallis, Don. 2000. "The Reliability of Randomized Algorithms." *British Journal for the Philosophy of Science* 51:255-71.

Fallis, Don. 2002. "What Do Mathematicians Want?: Probabilistic Proofs and the Epistemic Goals of Mathematicians." *Logique et Analyse* 45:373-88.

Fallis, Don. 2003. "Intentional Gaps in Mathematical Proofs." *Synthese* 134:45-69.

Fallis, Don. 2006. "The Epistemic Costs and Benefits of Collaboration." *Southern Journal of Philosophy* 44 (Supplement):197-208.

Fallis, Don. 2007. "Collective Epistemic Goals." *Social Epistemology* 21:267-80.

Fricker, Elizabeth. 2006. "Testimony and Epistemic Autonomy." Pp. 225-50 in *The Epistemology of Testimony*, eds. Jennifer Lackey and Ernest Sosa. Oxford: Oxford.

Goldman, Alvin. 1987. "Foundations of Social Epistemics." *Synthese* 73:109-44.

Goldman, Alvin I. 1999. *Knowledge in a Social World*. New York: Oxford.

Hakli, Raul. 2006. "Group Beliefs and the Distinction Between Belief and Acceptance." *Cognitive Systems Research* 7:286-97.

Hardwig, John. 1991. "The Role of Trust in Knowledge." *Journal of Philosophy* 88:693-708.

Huemer, Michael. 2005. "Is Critical Thinking Epistemically Responsible?" *Metaphilosophy* 36:522-31.

Hume, David. 1967 [1739]. *A Treatise of Human Nature*. ed. L. A. Selby-Bigge. London: Oxford.

Jackson, Jeffrey C. 2009. "Randomized Arguments are Transferable." *Philosophia Mathematica* 17:363-68.

Kitcher, Philip. 1993. *The Advancement of Science*. New York: Oxford.

Maddy, Penelope. 1997. *Naturalism in Mathematics*. New York: Oxford.

Mathiesen, Kay. 2006. "Epistemic Risk and Community Policing." *Southern Journal of Philosophy* 44 (Supplement):139-50.

McEvoy, Mark. 2008. "The Epistemological Status of Computer-Assisted Proofs." *Philosophia Mathematica* 16:374-87.

Myers, David G. 2001. "Do We Fear the Right Things?" *American Psychological Society Observer*. http://www.davidmyers.org/fears/

Nasar, Sylvia and David Gruber. 2006. "Manifold Destiny." *New Yorker*. http://www.newyorker.com/archive/2006/08/28/060828fa_fact2

Peressini, Anthony. 2003. "Proof, Reliability, and Mathematical Knowledge." *Theoria* 69:211-32.

Rabin, M. O. 1980. "Probabilistic Algorithm for Testing Primality." *Journal of Number Theory* 12:128-38.

Tymoczko, Thomas. 1979. "The Four-Color Problem and Its Philosophical Significance." *Journal of Philosophy* 76:57-83.

Warnken, Byron and Elizabeth Samuels. 1997. "The Types of Opinions Issued." *Online Academic Support Program for Law Students*. http://www.onlineasp.org/class/case08.htm

Whitehead, Alfred N. and Bertrand Russell. 1962 [1910]. *Principia Mathematica to *56*. Cambridge: Cambridge.

Wolfram, Stephen. 2002. *A New Kind of Science*. Champaign, Illinois: Wolfram Media, Inc.

Zeilberger, Doron. 1993. "Theorems for a Price: Tomorrow's Semi-Rigorous Mathematical Culture." *Notices of the American Mathematical Society* 40:978-81.